UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P O Box 1450
Alexandria, Virginia 22313-1450
www uspto gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/608,037	06/30/2003	Sanjay Ghemawat	0026-0031	8255

44989 7590 08/19/2009

HARRITY & HARRITY, LLP
11350 Random Hills Road
SUITE 600
FAIRFAX, VA 22030

EXAMINER	
LOVEL, KIMBERLY M	

ART UNIT	PAPER NUMBER
2167	

MAIL DATE	DELIVERY MODE
08/19/2009	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No.	Applicant(s)	
	10/608,037	GHEMAWAT ET AL.	
	Examiner	Art Unit	
	KIMBERLY LOVEL	2167	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE *3* MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *06 May 2009*.

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1,3-11,13,15,16 and 19-39* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1, 3-11, 13, 15, 16 and 19-39* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)

2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)

3) ☐ Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date _____.

4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____ .

5) ☐ Notice of Informal Patent Application

6) ☐ Other: _____.

DETAILED ACTION

1. Claims 1, 3-11, 13, 15, 16 and 19-39 are rejected.

2. In view of the Appeal Brief filed on 6 May 2009, PROSECUTION IS HEREBY REOPENED.

To avoid abandonment of the application, appellant must exercise one of the following two options:

(1) file a reply under 37 CFR 1.111 (if this Office action is non-final) or a reply under 37 CFR 1.113 (if this Office action is final); or,

(2) initiate a new appeal by filing a notice of appeal under 37 CFR 41.31 followed by an appeal brief under 37 CFR 41.37. The previously paid notice of appeal fee and appeal brief fee can be applied to the new appeal. If, however, the appeal fees set forth in 37 CFR 41.20 have been increased since they were previously paid, then appellant must pay the difference between the increased fees and the amount previously paid.

A Supervisory Patent Examiner (SPE) has approved of reopening prosecution by signing below:

/John R. Cottingham/

Supervisory Patent Examiner, Art Unit 2167

Claim Rejections - 35 USC § 103

3. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the

invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

4. Claims 1, 3-10, 13, 15, 16, 19, 23, 24, 27, 30, 31 and 34-39 are rejected under 35 U.S.C. 103(a) as being unpatentable over US PGPub 20020120763 to Miloushev et al (hereafter Miloushev) in view of US PGPub 2002/0191311 to Ulrich et al (hereafter Ulrich) in view of US PGPub 2003/0115439 to Mahalingam et al (hereafter Mahalingam).

Referring to claim 1, Miloushev discloses a file system, comprising:

a plurality of servers [file servers 101-107] (see [0124] and Fig 1); and

a master [file switch 100] connected to the servers [file servers 101-107] (see Fig 1) and configured to:

store namespace data that includes file identifiers for files [namespace aggregation] (see [0170])

where the master is configured to:

communicate with the servers at startup of the master to identify files stored by the servers (see [0302]).

While Miloushev discloses the striping of files over a plurality of file servers (see [0118]), Miloushev fails to explicitly disclose that the file servers are stored as chunks. Furthermore, while Miloushev discloses a master [file switch] which maintains file metadata that can be defined as a data structure containing information about how the contents of a given file is stored on the file servers, Miloushev fails to explicitly disclose the further limitations of the a master configured to: store mapping data that maps the file identifiers to the chunks to which the file identifiers correspond, store an operation

log that includes a record of changes to at least one of the namespace data or the mapping data, and store location data that identifies which of the servers stores which of the chunks. Also, while Miloushev discloses the collection of configuration information by the file switch during startup of the file switch (see [0302]), Miloushev fails to explicitly disclose the further limitation of where the master is configured to: communicate with the servers at startup of the master to identify the chunks stored by the servers, and record, in a non-persistent manner, information regarding the chunks stored by each of the servers as the location data.

Ulrich discloses a distributed file system, including the further limitations of servers storing files as chunks [blocks] (see [0122]),

a master connected to the servers and configured to:

store namespace data that includes file identifiers for files for which the file data is stored as chunks [Filename Table 310] (see [0185]; Fig 3; and Fig 14A),

store mapping data that maps the file identifiers to the chunks to which the file identifiers correspond (see [0183]-[0188]; Fig 3; and Fig 14A),

store location data that identifies which of the servers stores which of the chunks [Gee Table 320] (see [0187]; Fig 3; and Fig 14A),

where the master is configured to:

communicate with the servers to identify the chunks stored by the servers, and record, in a non-persistent manner [volatile memory], information regarding the chunks stored by each of the servers as the location data (see [0196]).

It would have been obvious to one of ordinary skill in the art at the time of the invention to utilize the metadata disclosed by Ulrich, which stored in a non-persistent manner as additional metadata collected by the files system of Miloushev. One would have been motivated to do so since Miloushev stripes data across disks, and maintaining metadata of the stripes increases efficiency of the system.

While the combination of Miloushev and Ulrich (hereafter Miloushev/Ulrich) teaches an intent log (Ulrich: see [0039]), Miloushev/Ulrich fails to explicitly disclose the further limitation of an operation log that includes a record of changes to at least one of the namespace data or the mapping data. Mahalingam discloses migration in distributed file system, including the further limitation of an operation log that includes a record of changes to at least one of the namespace data or the mapping data [namespace update process is logged to the intention file] (see [0042] and Fig 2).

It would have been obvious to one of ordinary skill in the art at the time of the invention to store the information of Mahalingam in the log of Ulrich. One would have been motivated to do so in order to increase access speeds of locating data through maintenance of location data.

Referring to claim 3, the combination of Miloushev/Ulrich and Mahalingam (hereafter Miloushev/Ulrich/Mahalingam) discloses the system of claim 1, wherein where the master is further configured to control placement of new chunks at the servers (Miloushev: striping individual files onto a plurality of servers; when a new file is to be created, the file switch selects the target file server – see [0118]; [0182]; [0285]; and [0286]; Ulrich: see [0405]-[0407]).

Referring to claim 4, Miloushev/Ulrich/Mahalingam discloses the system of claim 3, wherein where when controlling the placement of new chunks, the master is configured to:

identify one or more of the servers to store the new chunks based on at least one of utilization of the servers, prior chunk distribution involving the servers, network topology, or failure correlation properties associated with the servers (Miloushev: utilization of the storage capacity of the file servers and their relative load – see [0383]; [0384]; and [0388]; Ulrich: server utilization statistics – see [0137] and [0140]), and

place the new chunks at the identified one or more servers (Miloushev: switches the file create transaction to that server – see [0182]; Ulrich: the blocks are stored in the disk array – see [0410]).

Referring to claim 5, Miloushev/Ulrich/Mahalingam discloses the system of claim 1, where the master is further configured to control redistribution of the chunks stored by the servers (Miloushev: a subsystem in the file switch may determine a different way in which a given file or set of files should be distributed on the file servers – see [0384]; Ulrich: see [0460] – block redistribution).

Referring to claim 6, Miloushev/Ulrich/Mahalingam discloses the system of claim 5, wherein where when controlling redistribution of the chunks, the master is configured to:

select a chunk to redistribute based on a current distribution of the chunks,

identify one or more of the servers to which to move the selected chunk, and

move the selected chunk to the identified one or more servers [determine a different way in which a file or set of files should be distributed on the file servers] (Miloushev: determine a different way in which a file or set of files should be distributed on the file servers – see [0383]-[0388]; Ulrich: see [0397]-[0404] and [0425]).

Referring to claim 7, Miloushev/Ulrich/Mahalingam discloses the system of claim 1, wherein where the master is further configured to monitor a state of the servers (Miloushev: heartbeat that the file switch periodically sends to the server; state data for each file server – see [0191] and [0277]; Ulrich: see [0183]-[0188]).

Referring to claim 8, Miloushev/Ulrich/Mahalingam discloses the system of claim 7, wherein where the master [file switch] is configured to exchange heartbeat signals [heartbeat that the file switch periodically sends to the server] with the servers to determine the state of the servers [state data for each file server] (Miloushev: see [0191] and [0277]).

Referring to claim 9, Miloushev/Ulrich/Mahalingam discloses the system of claim 8, wherein where the heartbeat signals include space utilization information (Miloushev: see [0388]; Ulrich: see [0140] and [0507]).

Referring to claim 10, Miloushev/Ulrich/Mahalingam discloses the system of claim 7, wherein where the state of the servers includes information regarding the chunks [blocks] stored by the servers (Ulrich: see [0183]-[0188]).

Referring to claim 13, Miloushev discloses a master [file switch 100] in a file system that includes the master connected to a plurality of servers [file servers 101-107] (see [0124] and Fig 1), the master comprising:

means for communicating with the servers at startup of the master to identify files stored by the servers (see [0302]).

While Miloushev discloses the striping of files over a plurality of file servers (see [0118]), Miloushev fails to explicitly disclose that the file servers are stored as chunks. Furthermore, while Miloushev discloses a master [file switch] which maintains file metadata that can be defined as a data structure containing information about how the contents of a given file is stored on the file servers, Miloushev fails to explicitly disclose the further limitations of the a master configured to: store mapping data that maps the file identifiers to the chunks to which the file identifiers correspond, store an operation log that includes a record of changes to at least one of the namespace data or the mapping data, and store location data that identifies which of the servers stores which of the chunks. Also, while Miloushev discloses the collection of configuration information by the file switch during startup of the file switch (see [0302]), Miloushev fails to explicitly disclose the further limitation of means for storing, in a non-persistent manner, location information regarding the chunks stored by each of the servers.

Ulrich discloses a distributed file system, including the further limitations of servers storing files as chunks [blocks] (see [0122]),

means for storing namespace data that includes file identifiers for files for which the file data is stored as chunks [Filename Table 310] (see [0185]; Fig 3; and Fig 14A),

means for storing mapping data that maps the file identifiers to the chunks to which the file identifiers correspond (see [0183]-[0188]; Fig 3; and Fig 14A),

means for storing location data that identifies which of the servers stores which of the chunks [Gee Table 320] (see [0187]; Fig 3; and Fig 14A),

and means for updating location information by periodically instructing the servers to identify the data stored by the servers (see [0196]).

It would have been obvious to one of ordinary skill in the art at the time of the invention to utilize the metadata disclosed by Ulrich, which stored in a non-persistent manner as additional metadata collected by the files system of Miloushev. One would have been motivated to do so since Miloushev stripes data across disks, and maintaining metadata of the stripes increases efficiency of the system.

While the combination of Miloushev and Ulrich (hereafter Miloushev/Ulrich) teaches an intent log (Ulrich: see [0039]), Miloushev/Ulrich fails to explicitly disclose the further limitation of an operation log that includes a record of changes to at least one of the namespace data or the mapping data. Mahalingam discloses migration in distributed file system, including the further limitation of an operation log that includes a record of changes to at least one of the namespace data or the mapping data [namespace update process is logged to the intention file] (see [0042] and Fig 2).

It would have been obvious to one of ordinary skill in the art at the time of the invention to store the information of Mahalingam in the log of Ulrich. One would have been motivated to do so in order to increase access speeds of locating data through maintenance of location data.

Referring to claim 15, Miloushev discloses a file system, comprising:

a plurality of servers [file servers 101-107] (see [0124] and Fig 1); and

a master [file switch 100] connected to the servers [file servers 101-107] (see Fig 1) and configured to:

store namespace data that includes file identifiers for files [namespace aggregation] (see [0170])

where the master is configured to:

communicate with the servers at startup of the master to identify files stored by the servers (see [0302]).

While Miloushev discloses the striping of files over a plurality of file servers (see [0118]), Miloushev fails to explicitly disclose that the file servers are stored as chunks. Furthermore, while Miloushev discloses a master [file switch] which maintains file metadata that can be defined as a data structure containing information about how the contents of a given file is stored on the file servers, Miloushev fails to explicitly disclose the further limitations of the a master configured to: store mapping data that maps the file identifiers to the chunks to which the file identifiers correspond, store an operation log that includes a record of changes to at least one of the namespace data or the mapping data, and store location data that identifies which of the servers stores which of the chunks. Also, while Miloushev discloses the collection of configuration information by the file switch during startup of the file switch (see [0302]), Miloushev fails to explicitly disclose the further limitation of where the master is configured to: determine location information by communicating with the servers, the location information being based on which of the servers store ones of the chunks, store the

location information as the location data, and update the location data by periodically communicating with the servers to obtain changes to the location data.

Ulrich discloses a distributed file system, including the further limitations of servers storing files as chunks [blocks] (see [0122]),

a master connected to the servers and configured to:

store namespace data that includes file identifiers for files for which the file data is stored as chunks [Filename Table 310] (see [0185]; Fig 3; and Fig 14A),

store mapping data that maps the file identifiers to the chunks to which the file identifiers correspond (see [0183]-[0188]; Fig 3; and Fig 14A),

store location data that identifies which of the servers stores which of the chunks [Gee Table 320] (see [0187]; Fig 3; and Fig 14A),

where the master is configured to:

determine location information by communicating with the servers, the location information being based on which of the servers store ones of the chunks, store the location information as the location data, and update the location data by periodically communicating with the servers to obtain changes to the location data (see [0196]).

It would have been obvious to one of ordinary skill in the art at the time of the invention to utilize the metadata disclosed by Ulrich, which stored in a non-persistent manner as additional metadata collected by the files system of Miloushev. One would have been motivated to do so since Miloushev stripes data across disks, and maintaining metadata of the stripes increases efficiency of the system.

While the combination of Miloushev and Ulrich (hereafter Miloushev/Ulrich) teaches an intent log (Ulrich: see [0039]), Miloushev/Ulrich fails to explicitly disclose the further limitation of an operation log that includes a record of changes to at least one of the namespace data or the mapping data. Mahalingam discloses migration in distributed file system, including the further limitation of an operation log that includes a record of changes to at least one of the namespace data or the mapping data [namespace update process is logged to the intention file] (see [0042] and Fig 2).

It would have been obvious to one of ordinary skill in the art at the time of the invention to store the information of Mahalingam in the log of Ulrich. One would have been motivated to do so in order to increase access speeds of locating data through maintenance of location data.

Referring to claim 16, Miloushev discloses a file system, comprising:

a plurality of servers [file servers 101-107] (see [0124] and Fig 1); and

a master [file switch 100] connected to the servers [file servers 101-107] (see Fig 1) and configured to:

store namespace data that includes file identifiers for files [namespace aggregation] (see [0170])

where the master is configured to:

communicate with the servers at startup of the master to identify files stored by the servers (see [0302]).

While Miloushev discloses the striping of files over a plurality of file servers (see [0118]), Miloushev fails to explicitly disclose that the file servers are stored as chunks.

Furthermore, while Miloushev discloses a master [file switch] which maintains file metadata that can be defined as a data structure containing information about how the contents of a given file is stored on the file servers, Miloushev fails to explicitly disclose the further limitations of the a master configured to: store mapping data that maps the file identifiers to the chunks to which the file identifiers correspond, store an operation log that includes a record of changes to at least one of the namespace data or the mapping data, and store location data that identifies which of the servers stores which of the chunks. Also, while Miloushev discloses the collection of configuration information by the file switch during startup of the file switch (see [0302]), Miloushev fails to explicitly disclose the further limitation of where the master is configured to: communicate with the servers to determine location information of the data, the location information being based on which of the servers store the chunks, and store the location information as the location data.

Ulrich discloses a distributed file system, including the further limitations of servers storing files as chunks [blocks] (see [0122]),

a master connected to the servers and configured to:

store namespace data that includes file identifiers for files for which the file data is stored as chunks [Filename Table 310] (see [0185]; Fig 3; and Fig 14A),

store mapping data that maps the file identifiers to the chunks to which the file identifiers correspond (see [0183]-[0188]; Fig 3; and Fig 14A),

store location data that identifies which of the servers stores which of the chunks [Gee Table 320] (see [0187]; Fig 3; and Fig 14A),

where the master is configured to:

communicate with the servers to determine location information of the data, the location information being based on which of the servers store the chunks, and store the location information as the location data (see [0196]).

It would have been obvious to one of ordinary skill in the art at the time of the invention to utilize the metadata disclosed by Ulrich, which stored in a non-persistent manner as additional metadata collected by the files system of Miloushev. One would have been motivated to do so since Miloushev stripes data across disks, and maintaining metadata of the stripes increases efficiency of the system.

While the combination of Miloushev and Ulrich (hereafter Miloushev/Ulrich) teaches an intent log (Ulrich: see [0039]), Miloushev/Ulrich fails to explicitly disclose the further limitation of an operation log that includes a record of changes to at least one of the namespace data or the mapping data. Mahalingam discloses migration in distributed file system, including the further limitation of an operation log that includes a record of changes to at least one of the namespace data or the mapping data [namespace update process is logged to the intention file] (see [0042] and Fig 2).

It would have been obvious to one of ordinary skill in the art at the time of the invention to store the information of Mahalingam in the log of Ulrich. One would have been motivated to do so in order to increase access speeds of locating data through maintenance of location data.

Referring to claim 19, Miloushev/Ulrich/Mahalingam discloses the file system of claim 1, where the file identifiers are organized hierarchically in a tree of directories (Miloushev: see [0179] and [0389]; Ulrich: see [0184]).

Referring to claim 23, Miloushev/Ulrich/Mahalingam discloses the file system of claim 1, where the master is configured to update the location data by periodically instructing the servers to provide information [polling] regarding the chunks stored by the servers (Ulrich: see [0113]; [0131]; and [0183]-[0187]).

Referring to claim 24, Miloushev/Ulrich/Mahalingam discloses the file system of claim 1, where the operation log includes a logical timeline that defines an order for concurrent operations (Mahalingam: see [0035]).

Referring to claim 27, Miloushev/Ulrich/Mahalingam discloses the master of claim 13, where the operation log includes a logical timeline that defines an order for concurrent operations (Mahalingam: see [0035]).

Referring to claim 30, Miloushev discloses a method performed by a master device [file switch 100] in a file system that includes the master device connected to a plurality of server devices [file servers 101-107] (see [0124] and Fig 1), the method comprising:

communicating with the server devices to identify files stored by the servers (see [0302]); and

storing namespace data that includes file identifiers for files [namespace aggregation] (see [0170]).

While Miloushev discloses the striping of files over a plurality of file servers (see [0118]), Miloushev fails to explicitly disclose that the file servers are stored as chunks. Furthermore, while Miloushev discloses a master [file switch] which maintains file metadata that can be defined as a data structure containing information about how the contents of a given file is stored on the file servers, Miloushev fails to explicitly disclose the further limitations of storing location data that identifies which of the servers stores which of the chunks; storing mapping data that maps the file identifiers to the chunks to which the file identifiers correspond, maintaining an operation log that includes a record of changes to at least one of the namespace data and the mapping data. Ulrich discloses a distributed file system, including the further limitations of servers storing files as chunks [blocks] (see [0122]),

storing namespace data that includes file identifiers for files for which the file data is stored as chunks [Filename Table 310] (see [0185]; Fig 3; and Fig 14A),

storing mapping data that maps the file identifiers to the chunks to which the file identifiers correspond (see [0183]-[0188]; Fig 3; and Fig 14A),

storing location data that identifies which of the servers stores which of the chunks [Gee Table 320] (see [0187]; Fig 3; and Fig 14A),

communicating with the servers to identify the chunks stored by the servers (see [0196]).

It would have been obvious to one of ordinary skill in the art at the time of the invention to utilize the metadata disclosed by Ulrich, which stored in a non-persistent manner as additional metadata collected by the files system of Miloushev. One would

have been motivated to do so since Miloushev stripes data across disks, and maintaining metadata of the stripes increases efficiency of the system.

While the combination of Miloushev and Ulrich (hereafter Miloushev/Ulrich) teaches an intent log (Ulrich: see [0039]), Miloushev/Ulrich fails to explicitly disclose the further limitation of an operation log that includes a record of changes to at least one of the namespace data or the mapping data. Mahalingam discloses migration in distributed file system, including the further limitation of an operation log that includes a record of changes to at least one of the namespace data and the mapping data [namespace update process is logged to the intention file] (see [0042] and Fig 2).

It would have been obvious to one of ordinary skill in the art at the time of the invention to store the information of Mahalingam in the log of Ulrich. One would have been motivated to do so in order to increase access speeds of locating data through maintenance of location data.

Referring to claim 31, Miloushev/Ulrich/Mahalingam discloses the method of claim 30, where the operation log includes a logical timeline that defines an order for concurrent operations (Mahalingam: see [0035]).

Referring to claim 34, Miloushev/Ulrich/Mahalingam discloses the file system of claim 15, where the master is further configured to: identify one or more of the servers to store a new chunk based on prior chunk distribution involving the servers, and place the new chunk at the identified one or more servers (Miloushev: see [0190]; Ulrich: see Fig 24A).

Referring to claim 35, Miloushev/Ulrich/Mahalingam discloses the file system of claim 15, where the master is further configured to: identify one or more of the servers to store a new chunk based on failure correlation properties associated with the servers, and place the new chunk at the identified one or more servers (Miloushev: see [0190]; Ulrich: see Fig 24A and [0507]).

Referring to claim 36, Miloushev/Ulrich/Mahalingam discloses the file system of claim 16, where the master is further configured to: identify one or more of the servers to store a new chunk based on prior chunk distribution involving the servers, and place the new chunk at the identified one or more servers (Miloushev: see [0190]; Ulrich: see Fig 24A).

Referring to claim 37, Miloushev/Ulrich/Mahalingam discloses the file system of claim 16, where the master is further configured to: identify one or more of the servers to store a new chunk based on failure correlation properties associated with the servers, and place the new chunk at the identified one or more servers (Miloushev: see [0190]; Ulrich: see Fig 24A and [0507]).

Referring to claim 38, Miloushev/Ulrich/Mahalingam discloses the master of claim 13, further comprising: means for identifying one or more of the servers to store a new chunk based on utilization of the servers, prior chunk distribution involving the servers, and failure correlation properties associated with the servers; and means for placing the new chunk at the identified one or more servers (Miloushev: see [0190]; Ulrich: see Fig 24A and [0507]).

Referring to claim 39, Miloushev/Ulrich/Mahalingam discloses the method of claim 30, further comprising: identifying one or more of the server devices to store a new chunk based on utilization of the server devices, prior chunk distribution involving the server devices, and failure correlation properties associated with the server devices; and placing the new chunk at the identified one or more server devices (Miloushev: see [0190]; Ulrich: see Fig 24A and [0507]).

5. Claims 11, 21 and 22 are rejected under 35 U.S.C. 103(a) as being unpatentable over US PGPub 20020120763 to Miloushev et al in view of US PGPub 2002/0191311 to Ulrich et al in view of US PGPub 2003/0115439 to Mahalingam et al as applied to claims 1 and 10 above above, and further in view of US Patent No 5,644,751 to Burnett (hereafter Burnett).

Referring to claim 11, Miloushev/Ulrich/Mahalingam fail to explicitly disclose the further limitation of where the information of claim 10 includes version numbers of the chunks. Burnett discloses a distributed file system, including the further limitation of where the information includes version numbers of the chunks (see column 8, lines 35-54).

It would have been obvious at the time of the invention to utilize the information disclosed by Burnett with the system of Miloushev/Ulrich/Mahalingam. One would have been motivated to do so in order to increase the efficiency of the system by maintaining version information.

Referring to claim 21, Miloushev/Ulrich/Mahalingam fail to explicitly disclose the further limitation of where the master is configured to identify one of the chunks via a chunk handle that uniquely identifies the one of the chunks. Burnett discloses a distributed file system, including the further limitation of where the master is configured to identify one of the chunks via a chunk handle that uniquely identifies the one of the chunks (see column 8, lines 35-54).

It would have been obvious at the time of the invention to utilize the information disclosed by Burnett with the system of Miloushev/Ulrich/Mahalingam. One would have been motivated to do so in order to increase the efficiency of the system by maintaining version information.

Referring to claim 22, the combination of Miloushev/Ulrich/Mahalingam and Burnett discloses the file system of claim 21, where the chunk handle encodes a timestamp (Burnett: see column 8, lines 35-54).

6. Claim 20 is rejected under 35 U.S.C. 103(a) as being unpatentable over US PGPub 20020120763 to Miloushev et al in view of US PGPub 2002/0191311 to Ulrich et al in view of US PGPub 2003/0115439 to Mahalingam et al as applied to claim 1 above, and further in view of US Patent No 5,283,894 to Deran (hereafter Deran).

Referring to claim 20, while Miloushev/Ulrich/Mahalingam discloses a tree of directories, Miloushev/Ulrich/Mahalingam fails to explicitly disclose the further limitation of where the master stores the namespace data using prefix-compression. Deran

discloses the further limitation of applying prefix-compression to a tree (see column 22, lines 27-41).

It would have been obvious to one of ordinary skill in the art at the time of the invention to utilize the prefix-compression algorithm disclosed by Deran with the disclosed directory tree. One would have been motivated to do so since a directory can include a plurality of entries and prefix compression allows for increasing access speed of the directory.

7. Claims 25, 26, 28, 29, 32 and 33 are rejected under 35 U.S.C. 103(a) as being unpatentable over US PGPub 20020120763 to Miloushev et al in view of US PGPub 2002/0191311 to Ulrich et al in view of US PGPub 2003/0115439 to Mahalingam et al as applied to claim 1 above, and further in view of US Patent No 6,021,408 to Burnett (hereafter Burnett).

Referring to claims 25, 28 and 32, while Miloushev/Ulrich/Mahalingam discloses a log, fails to explicitly disclose the further limitation of where the master is configured to: determine when a size of the operation log exceeds a threshold, and create a checkpoint of the operation log when the size of the operation log exceeds the threshold. Ledain et al discloses migration in a distributed system, including the further limitation of determine when a size of the operation log exceeds a threshold, and create a checkpoint of the operation log when the size of the operation log exceeds the threshold (see column 9, line 62 – column 10, line 34; column 21, lines 22-36; column 25, line 57 – column 26, line 9; and column 29, line 58 – column 30, line 6).

It would have been obvious to one of ordinary skill in the art at the time of the invention to apply the steps disclosed by Ledain for maintaining a log to the log of Miloushev/Ulrich/Mahalingam. One would have been motivated to do so in order to decrease the recovery time by decreasing the size of the log.

Referring to claims 26, 29 and 33, the combination of Miloushev/Ulrich/Mahalingam and Ledain discloses the file system of claim 25, where the master is configured to: create a new operation log file, and create the checkpoint as a background operation (Ledain: see column 9, line 62 – column 10, line 34; column 21, lines 22-36; column 25, line 57 – column 26, line 9; and column 29, line 58 – column 30, line 6).

Response to Arguments

8. Applicant's arguments with respect to the claims have been considered but are moot in view of the new ground(s) of rejection.

Conclusion

Any inquiry concerning this communication or earlier communications from the examiner should be directed to KIMBERLY LOVEL whose telephone number is (571)272-2750. The examiner can normally be reached on 8:00 - 4:00.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, John Cottingham can be reached on (571) 272-7079. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/John R. Cottingham/
Supervisory Patent Examiner, Art Unit 2167

/Kimberly Lovel/
Examiner
Art Unit 2167

14 August 2009
/KL/